arXiv:1611.01542v1 [math.CO] 4 Nov 2016

THE SET SPLITTABILITY PROBLEM

PETER BERNSTEIN, CASHOUS BORTNER, SAMUEL COSKEY, SHUNI LI,
AND CONNOR SIMPSON

Abstract. The set splittability problem is the following: given a finite collection of finite sets, does there exits a single set that selects half the elements from each set in the collection? (If a set has odd size, we allow the floor or ceiling.) It is natural to study the set splittability problem in the context of combinatorial discrepancy theory and its applications, since a collection is splittable if and only if it has discrepancy ≤ 1.

After introducing the concepts and their background, we show that the set splittability problem is NP-complete. We in fact establish this for the generalized version called the p-splittability problem, in which one seeks to select the fraction p from each set instead of half. Next we investigate several criteria for splittability and p-splittability, giving a complete characterization of p-splittability for three sets and of splittability for four sets. Finally we show that when there are sufficiently many elements, unsplittability is asymptotically much more rare than splittability.

§1. Introduction

Let $\mathcal{B} = \{B_1, \ldots, B_n\}$ be a collection of finite sets and let $0 \leq p \leq 1$. We say the collection $\mathcal{B}$ is *p-splittable* if there exists a set S (called a *p-splitter*) such that for all $i \leq n$, we have that $|S \cap B_i| = \lfloor p|B_i| \rceil$, the nearest integer to $p|B_i|$. Of course, the nearest integer $\lfloor p|B_i| \rceil$ is not well-defined when $p|B_i|$ is a half-integer, and in this case we adopt the convention that $|S \cap B_i|$ may be either the floor or the ceiling. When $p = \frac{1}{2}$, we will sometimes simply say that $\mathcal{B}$ is *splittable* and that S is a *splitter*.

It is natural to study splittability and its generalizations in the context of combinatorial discrepancy theory. Given a collection $\mathcal{B}$ as above, the *discrepancy* of $\mathcal{B}$ is

$$\operatorname{disc}(\mathcal{B}) = \min_S \max_{i \leq n} \Big| |B_i \cap S| - |B_i \setminus S| \Big|.$$

Intuitively, the discrepancy measures to what extent it is possible to simultaneously and evenly split each set in the collection. In fact $\operatorname{disc}(\mathcal{B}) \leq 1$ if and only if $\mathcal{B}$ is ($p = \frac{1}{2}$) splittable.

2010 *Mathematics Subject Classification.* 05D05, 05C15, 11K38, 68Q17.

Upper bounds for the discrepancy have been studied in recent decades. In 1981 Beck and Fiala showed that if every element of $\bigcup B$ is contained in at most t of the sets in $\mathcal{B}$, then $\text{disc}(\mathcal{B}) \leq 2t - 2$ [BF81]. Incremental improvements to this bound can be found in works such as [BH97, Hel99, Buk13]. In 1985, Spencer gave an upper bound for the discrepancy of an arbitrary collection:

$$\text{disc}(\mathcal{B}) \leq K\sqrt{n}$$

where K is an absolute constant and n is the number of sets [Spe85]. The upper bound of $O(\sqrt{n})$ is asymptotically tight for general collections. Of course the discrepancy of any given collection may be much smaller than this bound, and since least discrepancy is usually best, it is natural to study the discrepancy ≤ 1 case.

Set splittability can also be viewed as a combinatorial version of the outcome of the ham sandwich theorem: given Lebesgue measurable subsets $B_1, \ldots, B_n \subseteq \mathbb{R}^n$ there exists a hyperplane H such that for all $i \leq n$ exactly half the measure of B_i lies to each side of H. If Lebesgue measure is replaced by a discontinuous measure, then some of the mass of B_i may lie on H itself. In this case the conclusion must be modified to say that that *at most* half the measure of B_i lies to each side of H [EH]. Thus a ham sandwich hyperplane does not precisely solve the splittability problem, nor does set splittability help to find a geometric hyperplane. Nevertheless the two problems are closely related.

A third way to think of set splittability is as a very strong form of hypergraph 2-colorability. Recall that a hypergraph with hyperedges $B_1, \ldots, B_n$ is 2-colorable if there exists a $\{\text{red}, \text{blue}\}$-coloring of its vertices such that no hyperedge is monochromatic. With p-splittability we ask not simply that both colors are represented in each hyperedge, but that the color red always appears a prescribed percentage of the time.

In the next section, we study the computational complexity of the set splittability problem. In particular, we will show that for any $0 < p < 1$, the p-splittability problem is NP-complete. Questions of computational complexity have of course been explored in both the discrepancy theory and the hypergraph coloring contexts. For example, the problem of deciding whether a given hypergraph is 2-colorable has been shown to be NP-complete [Lov73].

In the area of discrepancy theory, when Spencer established the above-mentioned bound on the discrepancy he also conjectured that it is not efficient to find a splitter S which actually *witnesses* the bound [AS00]. The conjecture was disproved in 2010, when such an algorithm was indeed found [Ban10]. In 2012, an efficient algorithm was found with the additional property that its correctness does not depend on Spencer's theorem itself, thus providing a constructive proof of Spencer's result [LM12]. In the context of these

positive results, it is natural to ask whether it is efficient to decide whether a collection has discrepancy ≤ 1. Our NP-completeness result, invoked in the special case $p = \frac{1}{2}$, answers this question in the negative. This result thus affirms the belief that some discrepancy decision problems are algorithmically hard.

The fact that the p-splittability decision problem is hard means we cannot hope to find a general and useful characterization of p-splittability. However it is possible to do so for small collections and for other special collections of sets. For an example involving (very) small collections, we will show that a collection $\mathcal{B}$ of at most two sets is p-splittable for any p. For an example involving special collections, suppose that $\mathcal{B} = \{B_1, \ldots, B_n\}$ is a collection of n sets such that every element x lies in exactly $n-1$ sets of $\mathcal{B}$. In this case we will show that $\mathcal{B}$ is p-splittable if and only if the sum $\lfloor p|B_1| \rceil + \cdots + \lfloor p|B_n| \rceil$ is divisible by $n-1$. The calculations used in these two results eventually lead us to a complete algebraic characterization of p-splittability for collections $\mathcal{B}$ of at most three sets.

If one specializes to the important case $p = \frac{1}{2}$, some things become simpler and new characterizations become tractable. For example, if $\mathcal{B}$ is a collection of three sets then $\mathcal{B}$ is $\frac{1}{2}$-unsplittable if and only if each Venn region of multiplicity 2 is odd, and all other Venn regions are empty. (This result was previously observed in [CCSS16].) With the help of a supercomputer, we were also able to provide a complete characterization of the $\frac{1}{2}$-unsplittable collections of four sets in terms of the sizes of its Venn regions. The proof of the result rests on an exhaustive search for unsplittable configurations with a small number of elements, together with a reduction lemma which implies that if $\mathcal{B}$ is unsplittable then it remains unsplittable after reducing the number of elements of each Venn region modulo 2.

These Venn-style characterizations of the $\frac{1}{2}$-unsplittable configurations easily imply that unsplittability is extremely rare for collections of three and four sets. Although our method of finding unsplittable configurations becomes intractable for collections of five or more sets, this rarity phenomenon remains true. Specifically we show that if $n \to \infty$ and k grows sufficiently fast relative to n, then the probability that a collection with n sets and k elements is splittable converges to 1. In particular if n is fixed and k is large enough then most collections with n sets and k elements are splittable.

The rest of this paper is organized as follows. In Section 2 we prove that the problem of deciding whether a given collection is p-splittable is NP-complete. In Section 3, we give criteria for deciding whether some special collections are p-splittable, and provide a complete characterization of p-splittability for collections of at most three sets. In Section 4, we give further splittability criteria for the special case $p = \frac{1}{2}$ and use them to give a

complete characterization of $\frac{1}{2}$-splittability for collections of at most four sets. Finally we show that for collections with sufficiently many elements, splittability is by far more common than unsplittability.

Acknowledgement. This article represents a portion of the research carried out during the 2016 math REU program at Boise State University. The program was supported by NSF grant #DMS 1359425 and by Boise State University.

§2. The complexity of splittability

In this section we establish that the set splittability problem is hard. Before proceeding, let us clarify how we regard the p-splittability problem formally as a decision problem, which we denote p-Split. If $\mathcal{B} = \{B_1, \ldots, B_n\}$ is a finite collection of subsets of $\{1, \ldots, m\}$ then the *incidence matrix* of $\mathcal{B}$ is the $n \times m$ matrix M whose (i, j) entry is 1 whenever $j \in B_i$, and 0 otherwise. An instance of p-Split consists of an incidence matrix M of a collection $\mathcal{B}$. The matrix M lies in p-Split if and only if there exists a binary vector $\mathbf{y}$ such that $(M\mathbf{y})_i = \lfloor p|B_i| \rceil$. Indeed, such a vector $\mathbf{y}$ is the indicator function of a set S that is a p-splitter of $\mathcal{B}$.

In this section we will make significant use of the notations: M_i for the ith row of M; $\mathbf{1}$ for a vector of 1's of length determined by context, and; $M_i\mathbf{1}$ for the number of 1's in the ith row of M. It is important to note that the values of the right-hand sides $\lfloor p|B_i| \rceil$ from the previous paragraph can be determined just from the matrix M because we have $\lfloor p|B_i| \rceil = \lfloor pM_i\mathbf{1} \rceil$.

Theorem 2.1. *For any $0 < p < 1$, the problem p-Split of determining whether a collection is p-splittable is NP-complete.*

Note that p-Split lies in NP because given an instance M of p-Split and a characteristic vector $\mathbf{y}$ of an ostensible splitter, one can easily decide in polynomial time whether $(M\mathbf{y})_i$ is equal to $\lfloor pM_i\mathbf{1} \rceil$ for each i.

To establish that p-Split is NP-complete, we will exhibit a polynomial-time reduction from the decision problem ZOE (which is very similar to zero-one integer programming [Kar72]) to p-Split. Here ZOE stands for *zero one equations*, and is formalized as follows. An instance of ZOE consists of a binary matrix A. The matrix A lies in ZOE if and only if there exists a binary vector $\mathbf{x}$ such that $A\mathbf{x} = \mathbf{1}$. It is known that ZOE is NP-complete [DPV06].

In the definition of ZOE, we can clearly assume without loss of generality that the matrix A has no zero rows. We can also assume that at least one row of A has at least two 1's. Indeed, if A has just one 1 in each row then $A\mathbf{x} = \mathbf{1}$ is trivial to solve.

Now in order to establish Theorem 2.1, we will describe a mapping from binary matrices A to incidence matrices M, with the property that A lies in ZOE if and only if M lies in p-Split. In order to guarantee this, the matrix M that we construct will have the special properties:

(a) A is an upper-left submatrix of M;
(b) any solution $\mathbf{x}$ to $A\mathbf{x} = \mathbf{1}$ extends to a solution $\mathbf{y}$ to $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$; and
(c) any solution $\mathbf{y}$ to $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ restricts to a solution $\mathbf{x}$ to $A\mathbf{x} = \mathbf{1}$.

Having described our general approach, we now proceed with the details.

§2.1. Specification of the construction. Let A be a given $a_r \times a_c$ binary matrix, and assume that $A_i\mathbf{1} > 0$ for all i and that at least one $A_i\mathbf{1} > 1$. Additionally let $p \leq \frac{1}{2}$ be given. We will now construct a block matrix M of the form:

$$M := \begin{bmatrix} A & B & C \\ 0 & D & E \end{bmatrix}.$$

Before we describe the blocks B, C, D, and E, let us let $s := \max_i A_i\mathbf{1}$ and let T be indices of the $s-1$ columns following the columns of A. Next let $q = \frac{1-p}{p}$ and let F be the indices of the $\max\{\lceil qs \rceil - s + 1, \lceil q \rceil + 1\}$ many columns to the right of the columns indexed by T.

Now B is an $a_r \times |T|$ matrix whose ith row contains $A_i\mathbf{1} - 1$ many 1's, followed by all 0's. And C is an $a_r \times |F|$ matrix whose ith row contains $\lceil qA_i\mathbf{1} \rceil - A_i\mathbf{1} + 1$ many 1's, followed by all 0's.

The blocks D and E each consist of sublocks. For $i \leq |T|$, let D_i be the $\binom{|F|}{\lceil q \rceil} \times |T|$ matrix whose ith column consists of 1's, and all other columns consist of 0's. And let E_0 denote a $\binom{|F|}{\lceil q \rceil} \times |F|$ matrix whose rows consist of the indicator functions of the subsets of $\{1, \ldots, |F|\}$ of size $\lceil q \rceil$. Then we let

$$D = \begin{bmatrix} D_1 \\ \vdots \\ D_{|T|} \end{bmatrix}, \text{ and } E = \begin{bmatrix} E_0 \\ \vdots \\ E_0 \end{bmatrix}$$

Here there are $|T|$ many copies of E_0 in E.

It is easy to see that the dimensions of the matrix M are polynomial in the dimensions of the matrix A. (Recall here that p is a fixed parameter of the construction.) Hence the construction is a polynomial-time mapping.

§2.2. Example of the construction. Before proving that the construction satisfies our requirements, let us give an example. Suppose that $p = \frac{1}{3}$ and we are given the ZOE

system

$$A\mathbf{x} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Then we have $A_1\mathbf{1} = 2$, $A_2\mathbf{1} = 3$, $s = \max_i A_i\mathbf{1} = 3$, and $q = \frac{1-p}{p} = 2$. Thus T consists of the $s - 1 = 2$ columns $\{5, 6\}$, and F consists of the $\lceil qs \rceil - s + 1 = 4$ columns $\{7, 8, 9, 10\}$.

The block B is thus a 2×2 matrix with $A_1\mathbf{1} - 1 = 1$ many 1's in the first row and $A_2\mathbf{1} - 1 = 2$ many 1's in the second row. The block C is a 2×4 matrix with $\lceil qA_1\mathbf{1} \rceil - A_1\mathbf{1} + 1 = 3$ many 1's in the first row and $\lceil qA_2\mathbf{1} \rceil - A_2\mathbf{1} + 1 = 4$ many 1's in the second row.

Next, the blocks D_1 and D_2 are each $\binom{|F|}{\lceil q \rceil} \times |T|$ which comes to 6×2. Block D_1 is a column of 1's followed by a column of 0's, and block D_2 is a column of 0's followed by a column of 1's.

Finally, the block E_0 is $\binom{|F|}{\lceil q \rceil} \times |F|$ which comes to 6×4. The 6 rows of E_0 correspond to the 6 subsets of $\{1, \ldots, |F|\} = \{1, \ldots, 4\}$ of size $\lceil q \rceil = 2$. The full matrix is displayed in Figure 1.

$$M\mathbf{y} = \left[\begin{array}{cccc|cc|cccc} 1 & 1 & 0 & 0 & 1 & & 1 & 1 & 1 & \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \\ \hline & & & & 1 & & 1 & 1 & & \\ & & & & 1 & & 1 & & 1 & \\ & & & & 1 & & 1 & & & 1 \\ & & & & 1 & & & 1 & 1 & \\ & & & & 1 & & & 1 & & 1 \\ & & & & 1 & & & & 1 & 1 \\ & & & & & 1 & 1 & 1 & & \\ & & & & & 1 & 1 & & 1 & \\ & & & & & 1 & 1 & & & 1 \\ & & & & & 1 & & 1 & 1 & \\ & & & & & 1 & & 1 & & 1 \\ & & & & & 1 & & & 1 & 1 \end{array}\right] \mathbf{y} = \begin{bmatrix} 2 \\ 3 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

FIGURE 1. The constructed matrix M, and the corresponding equation of a splittability decision problem.

Having given the example construction, we briefly preview how the proof will play out in this case. The matrix M gives rise to the collection $\{B_1, \ldots, B_{14}\}$ of subsets of $\{1, 2, \ldots, 10\}$, where the characteristic vector of B_i is the ith row of M. The collection is splittable if and only if the linear system $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ has a binary solution. In our

example the values one the right-hand side are $pM_1\mathbf{1} = 2$, $pM_2\mathbf{1} = 3$, and $pM_i\mathbf{1} = 1$ for $i > 2$.

Note that if $A\mathbf{x} = \mathbf{1}$, then $\mathbf{x}$ extends to a solution of $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ by setting the variables of T to be 1 and the variables of F to be 0. Conversely, if $M\mathbf{y} = \lfloor p|B_i| \rceil$ then carefully inspecting the D and E blocks of M, the variables of T are forced to be 1 and the variables of F are forced to be 0. Since the rows of the block B have one fewer 1 than the rows of A (and the corresponding right-hand side), such a $\mathbf{y}$ restricts to a solution to $A\mathbf{x} = \mathbf{1}$.

§2.3. Proof of the main theorem. We now establish that the construction described above is indeed a reduction from ZOE to p-Split. We will assume throughout that $0 < p \leq \frac{1}{2}$, since it is clear that a collection is p-splittable if and only if it is $1-p$ splittable. To begin, we present a simple rounding calculation that will be used below.

Lemma 2.2. *Let $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ as before. Then for any $m \in \mathbb{N}$ we have*

$$\lfloor p(m + \lceil qm \rceil) \rceil = m.$$

Proof. Let $\varepsilon = \lceil qm \rceil - qm$. Then

$$\begin{aligned} \lfloor p(m + \lceil qm \rceil) \rceil &= \lfloor p(m + qm + \varepsilon) \rceil \\ &= \lfloor pm + m(1-p) + p\varepsilon \rceil \\ &= \lfloor m + p\varepsilon \rceil \\ &= m + \lfloor p\varepsilon \rceil \end{aligned}$$

Since $\varepsilon < 1$ and $p \leq \frac{1}{2}$, we have $p\varepsilon < \frac{1}{2}$, which gives that the last quantity equals m as desired. □

Next we establish the values of the right-hand sides of the system $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ that we have constructed.

Lemma 2.3. *Let $0 < p \leq \frac{1}{2}$, let A be $a_r \times a_c$, and let M be constructed as above. Then $\lfloor pM_i\mathbf{1} \rceil = A_i\mathbf{1}$ for all $i \leq a_r$, and $\lfloor pM_i\mathbf{1} \rceil = 1$ for all $i > a_r$.*

Proof. First consider $i \leq a_r$. Then

$$\begin{aligned} \lfloor pM_i\mathbf{1} \rceil &= \lfloor p(A_i\mathbf{1} + B_i\mathbf{1} + C_i\mathbf{1}) \rceil \\ &= \left\lfloor p\Big(A_i\mathbf{1} + (A_i\mathbf{1} - 1) + \big(\lceil qA_i\mathbf{1} \rceil - A_i\mathbf{1} + 1 \big) \Big) \right\rceil \\ &= \lfloor p(A_i\mathbf{1} + \lceil qA_i\mathbf{1} \rceil) \rceil \end{aligned}$$

By Lemma 2.2, the latter quantity is simply $A_i\mathbf{1}$, as claimed.

Next consider $i > a_r$. Here we have

$$\lfloor pM_i\mathbf{1} \rceil = \lfloor p(D_i\mathbf{1} + E_i\mathbf{1}) \rceil = \lfloor p(1 + \lceil q \rceil) \rceil .$$

Again using Lemma 2.2, the latter quantity is 1, as desired. □

To commence with the proof proper, we first show that if $A\mathbf{x} = \mathbf{1}$ has a solution, then $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ has a solution. Given a solution $\mathbf{x}$ to $A\mathbf{x} = \mathbf{1}$, we extend $\mathbf{x}$ to a vector $\mathbf{y}$ by appending $|T|$ many 1's followed by $|F|$ many 0's. Then for $i \leq a_r$ we have

$$M_i\mathbf{y} = A_i\mathbf{x} + B_i\mathbf{1} + C_i\mathbf{0} = 1 + (A_i\mathbf{1} - 1) = A_i\mathbf{1}$$

By Lemma 2.3, $A_i\mathbf{1} = \lfloor pM_i\mathbf{1} \rceil$. On the other hand for $i > a_r$, we have

$$M_i\mathbf{y} = 0\mathbf{x} + D_i\mathbf{1} + E\mathbf{0} = D_i\mathbf{1} = 1$$

Again by Lemma 2.3, this is equal to $\lfloor pM_i\mathbf{1} \rceil$, as desired.

For the converse, we show that if $(M\mathbf{y})_i = \lfloor pM_i\mathbf{1} \rceil$ has a solution then $A\mathbf{x} = \mathbf{1}$ has a solution. We make a series of claims about the structure of $\mathbf{y}$ that will enable us to create from it a solution $\mathbf{x}$ to $A\mathbf{x} = \mathbf{1}$.

Claim 2.4. Either there is an index $j \in T$ such that $y_j = 1$ or there is an index $k \in F$ such that $y_k = 1$, but not both.

Proof of claim. Suppose towards a contradiction that $j \in T$, $k \in F$ and $y_j = y_k = 1$. Recalling the definitions of D and E, we can find a row index $i > a_r$ such that M_i has a 1 in its jth and kth columns. It follows that $M_i\mathbf{y} \geq 2$, which contradicts the calculation from Lemma 2.3 that $M_i\mathbf{y} = \lfloor pM_i\mathbf{1} \rceil = 1$. ⊣

Claim 2.5. If $0 < p < \frac{1}{2}$, then for all indices $i \in T$, $y_i = 1$ and for all indices $j \in F$, $y_j = 0$.

Proof of claim. Suppose towards a contradiction that there is a $j \in F$ such $y_j = 1$. By the previous claim, for all $i \in T$ we have $y_i = 0$. If there is just one such $j \in F$ with $y_j = 1$, then by construction of E_0 we can find a row $\ell > a_r$ such that the jth entry of M_ℓ is 0. This implies that $M_\ell\mathbf{y} = 0$ contradicting that $M_\ell\mathbf{y} = \lfloor pM_\ell\mathbf{1} \rceil = 1$.

On the other hand if there are two distinct $j, j' \in F$ with $y_j = y_{j'} = 1$, then since $p < \frac{1}{2}$ implies $q > 1$, we can find a row $\ell > a_r$ such that the jth and j'th entry of M_ℓ is 1. This implies that $M_\ell\mathbf{y} \geq 2$ again contradicting that $M_\ell\mathbf{y} = 1$.

Thus we have shown that $y_j = 0$ for all $j \in F$. It follows from the construction of D that $y_i = 1$ for all $i \in T$. ⊣

To continue the proof, let us assume first that $p < \frac{1}{2}$. Then for all $i \in T$ we have $y_i = 1$ and for all $j \in F$ we have $y_j = 0$. Letting $\mathbf{x}$ denote the restriction of $\mathbf{y}$ to its first a_c entries, for any $i \leq a_r$ we have $(M\mathbf{y})_i = A\mathbf{x} + (A_i\mathbf{1} - 1)$. By Lemma 2.3 we also know that $(M\mathbf{y})_i = A_i\mathbf{1}$. It follows that $A\mathbf{x} = 1$.

Next consider the case when $p = \frac{1}{2}$. Then $q = 1$ so both D and E have exactly one 1 per row. It follows that we either have

I. $y_i = 1$ for all $i \in T$ and $y_j = 0$ for all $j \in F$, or
II. $y_i = 0$ for all $i \in T$ and $y_j = 1$ for all $j \in F$.

If I holds, we can do as we did when $p < \frac{1}{2}$, so we are done. Otherwise, if II holds, let $\mathbf{y}' = \mathbf{1} - \mathbf{y}$. Then

$$M_i\mathbf{y}' = M_i(\mathbf{1} - \mathbf{y}) = M_i\mathbf{1} - \lfloor M_i\mathbf{1}p \rceil .$$

We know that $A_i\mathbf{1}$ has the opposite parity of $A_i\mathbf{1} - 1 = B_i\mathbf{1}$, and that $C_i\mathbf{1} = 1$ when $p = \frac{1}{2}$. Therefore, $M_i\mathbf{1}$ is even, so $M_i\mathbf{1} - \lfloor M_i\mathbf{1}p \rceil = \lfloor M_i\mathbf{1}p \rceil$, meaning that $M_i\mathbf{y}' = \mathbf{b}$.

Thus, $\mathbf{y}'$ also corresponds to a valid splitter of $\mathcal{B}$, and since $y'_i = 1$ for all $i \in T$ we must also have that its first a_c entries pick out exactly one 1 per row of A by the argument used for when $p < \frac{1}{2}$. Therefore, taking $\mathbf{x}$ to be restriction of $\mathbf{y}'$ to its first a_c entries, we find that $A\mathbf{x} = \mathbf{1}$ once again.

This concludes the proof that the construction of M is a reduction from ZOE to p-Split.

§3. p-Splittability criteria and characterizations

The result of the previous section implies that it is not easy to find a general characterization of p-splittability. Nevertheless, in this section we provide several p-splittability criteria for special types of collections. Furthermore we completely characterize p-splittability for collections of at most three sets.

Before we begin our study, it is necessary to introduce the following notation. For a collection $\mathcal{B} = \{B_1, \ldots, B_n\}$ and an element x, the *multiplicity* of x is the number m_x of sets B_i such that $x \in B_i$. Given a subsequence of the sets $B_{i_1}, \ldots, B_{i_k}$ we can form a *Venn region* R of $\mathcal{B}$ consisting of the elements x that lie in $B_{i_1}, \ldots, B_{i_k}$ and in no other B_j. If R consists of elements of multiplicity k, we will also say that R has multiplicity k.

In the following result, we will say that a sequence $\{t_i\}$ is a *target sequence* for $\mathcal{B}$ if $0 \leq t_i \leq |B_i|$ for all i. The target sequence t_i is *achievable* if there is a set S such that $|S \cap B_i| = t_i$ for all i.

Lemma 3.1 (Parity Lemma)**.** *Let $\mathcal{B}$ be a collection and assume that for every $x \in \bigcup \mathcal{B}$ the multiplicity m_x is divisible by m. If the target sequence $\{t_i\}$ is achievable, then $\sum t_i$ is divisible by m.*

Proof. Let S be a set witnessing that $\{t_i\}$ is achievable. Then

$$\sum_{1 \le i \le n} t_i = \sum_{1 \le i \le n} |S \cap B_i| = \sum_{x \in S} m_x$$

By hypothesis, m_x is divisible by m for all x, so the right-hand side is also divisible by m. □

Since $\mathcal{B}$ is p-splittable if and only if the target sequence $t_i = \lfloor p|B_i| \rceil$ is achievable, the contrapositive of Lemma 3.1 provides a useful condition for showing that certain collections are not p-splittable. While the converse of Lemma 3.1 is false in general, we do have the following partial converse.

Lemma 3.2. *Let $\mathcal{B} = \{B_1, \ldots, B_n\}$ be a collection such that for all $x \in \bigcup \mathcal{B}$, $m_x = n - 1$. If $\sum \lfloor p|B_i| \rceil$ is divisible by $n - 1$, then $\mathcal{B}$ is p-splittable.*

Proof. Let $b_i = |\bigcup \mathcal{B} - B_i|$ be the size of the (unique) Venn region of multiplicity $n - 1$ which is not contained in B_i. Letting $t_i = \lfloor p|B_i| \rceil$ be the target sequence, we wish to find values $\bar{b}_i$ such that $0 \le \bar{b}_i \le b_i$ and $\sum_{j \ne i} \bar{b}_j = t_i$. The latter system of equations is square and invertible, so an elementary calculation gives the unique solution:

$$\bar{b}_i = \frac{1}{n-1} \left(-(n-2)t_i + \sum_{j \ne i} t_j \right). \tag{1}$$

Note that $\bar{b}_i$ is always an integer, because the above expression is equal to

$$\frac{1}{n-1} \left(-(n-1)t_i + \sum_{j} t_j \right) = -t_i + \frac{1}{n-1} \sum_{j} t_j$$

and $\sum t_j$ is divisible by $n - 1$ by hypothesis. Hence it remains only to establish that $0 \le \bar{b}_i \le b_i$.

For this, note that $t_i = \lfloor p|B_i| \rceil = \left\lfloor \sum_{j \ne i} p b_j \right\rceil = \epsilon_i + \sum_{j \ne i} p b_j$ where $|\epsilon_i| \le \frac{1}{2}$. Substituting this expression in for every t_i in Equation (1), we note that pb_i occurs $n - 1$ times in the parentheses while all other pb_j occur $n - 2$ times negatively and $n - 2$ times positively. Thus, the pb_j cancel, leaving us just with pb_i and error terms as follows:

$$\bar{b}_i = p b_i + \frac{1}{n-1} \left(-(n-2)\epsilon_i + \sum_{j \ne i} \epsilon_j \right). \tag{2}$$

There are $2n-3$ many ϵ terms in the parentheses, so we can conclude that $\bar{b}_i = pb_i + E$ where $|E| < 1$. Since $0 \leq pb_i \leq b_i$ and all of $0, b_i, \bar{b}_i$ are integers, it follows that $0 \leq \bar{b}_i \leq b_i$ too. □

Remark 3.3. It is necessary to amend the statement of the previous result if any of the $p|B_i|$ are half-integers. In this case the property "$\sum \lfloor p|B_i| \rceil$ is divisible by $n-1$" should be replaced by "it is possible to select the values of $\lfloor p|B_i| \rceil$ in such a way that $\sum \lfloor p|B_i| \rceil$ is divisible by $n-1$." Several of the results below will have similar amendments, as we will note briefly each time.

We are now ready to begin our classification of p-splittability for collections of size ≤ 3. We begin with the simple case of just two sets, because it helps motivate some of the steps for the three set case below.

Theorem 3.4. *Every collection of two sets is p-splittable for every $0 \leq p \leq 1$.*

Proof. Let $\mathcal{B} = \{B_1, B_2\}$ be a given two-set collection. Replacing p with $1-p$ if necessary, we can suppose that $p \leq \frac{1}{2}$. We fix the following notation for the sizes of the regions of $\mathcal{B}$: $a_1 = |B_1 \cap B_2^c|$, $a_2 = |B_1^c \cap B_2|$, and $b = |B_1 \cap B_2|$ (see Figure 2). Next let $t_1 = \lfloor p(a_1 + b) \rceil$ and $t_2 = \lfloor p(a_2 + b) \rceil$ denote the target cardinalities for $S \cap B_1$ and $S \cap B_2$ for a p-splitter S. To show $\mathcal{B}$ is p-splittable it suffices to find integers $\bar{a}_i, b$ such that: (i) $0 \leq \bar{a}_i \leq a_i$; (ii) $0 \leq \bar{b} \leq b$, and; (iii) $\bar{a}_i + \bar{b} = t_i$.

For this we let $\bar{b} = \lfloor pb \rceil$ and $\bar{a}_i = t_i - \bar{b}$ so that (ii) and (iii) are clearly satisfied. Of course the definitions of both t_i and $\bar{b}$ may be ambiguous; in such cases we need only make sure that if $a_i = 0$ then we choose $\bar{a}_i = 0$ too.

To see that (i) is satisfied, write $\varepsilon = \bar{b} - pb$ for the rounding error in computing $\bar{b}$ and $\epsilon_i = t_i - p(a_i + b)$ for the rounding error in computing t_i. Then the equations easily imply that we have

$$\bar{a}_i = pa_i + (\epsilon_i - \varepsilon).$$

Since $|\epsilon_i| \leq \frac{1}{2}$ and $|\varepsilon| \leq \frac{1}{2}$ we know that $|\epsilon_i - \varepsilon| \leq 1$. Assuming $a_i > 0$ the above equation gives $-1 < \bar{a}_i < a_i + 1$, and since $\bar{a}_i$ and a_i are integers, (i) is satisfied. On the other hand if $a_i = 0$ then $\bar{a}_i = 0$ too and (i) is clearly satisfied. □

To state our results for three sets, we extend the notation from the previous proof. For a three-set collection $\mathcal{B} = \{B_1, B_2, B_3\}$ we let a_i denote the number of multiplicity 1 elements of B_i, let b_i denote the number of multiplicity 2 elements *not* in B_i, and let c denote the number of multiplicity 3 elements (see Figure 2). As in the previous proof we

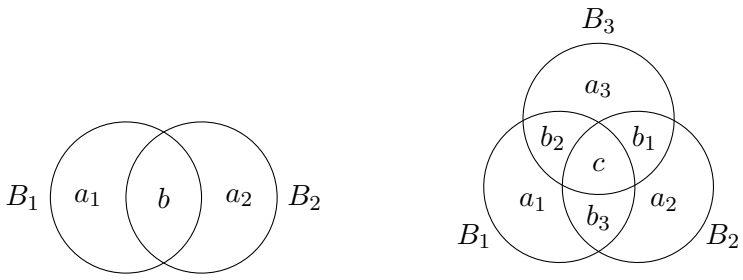

FIGURE 2. At left: sizes of the Venn regions of a two element collection $\{B_1, B_2\}$. At right: sizes of the Venn regions of a three element collection $\{B_1, B_2, B_3\}$.

let $t_i = \lfloor p|B_i| \rceil$ be the targets and $\epsilon_i = t_i - p|B_i|$ be the rounding error. Finally we set the values $\rho_i = -\epsilon_i + \sum_{j \neq i} \epsilon_j$.

Lemma 3.5. *Assume that $p \leq \frac{1}{2}$, and let $\mathcal{B} = \{B_1, B_2, B_3\}$ be given. Also assume there are no multiplicity 1 elements, that is, all $a_i = 0$. Then $\mathcal{B}$ is not p-splittable if and only if $\sum t_i$ is odd and one of the conditions holds:*

(a) *$c = 0$; or*
(b) *$pc < \frac{1}{2}$ and some $pb_i + \frac{1}{2}(pc - 1 + \rho_i) < 0$.*

Proof. Recall that $\mathcal{B}$ is p-splittable if and only if one can find values $\bar{b}_i$ and $\bar{c}$ such that $0 \leq \bar{b}_i \leq b$, $0 \leq \bar{c} \leq c$, and $\bar{c} + \sum_{j \neq i} \bar{b}_i = t_i$. Assuming one has chosen the value of $\bar{c}$, the equations have the unique solution:

$$\bar{b}_i = \frac{1}{2}\left(-t_i + \sum_{j \neq i} t_j - \bar{c}\right). \tag{3}$$

At this point we can observe that in order to achieve integer values of $\bar{b}_i$, one must choose the value $\bar{c}$ to have the same parity as $\sum t_j$. Next we substitute $t_j = pc + p\sum_{k \neq j} b_k + \epsilon_j$ to rewrite the above equation as

$$\bar{b}_i = pb_i + \frac{1}{2}\left(pc - \bar{c} + \rho_i\right). \tag{4}$$

Now assume that $\sum t_i$ is odd and that either (a) or (b) holds. We have already shown in Lemma 3.1 that if condition (a) holds then $\mathcal{B}$ is unsplittable. Next assume that (b) holds. Since $\sum t_i$ is odd, we cannot choose $\bar{c}$ to be of even parity and in particular cannot choose $\bar{c} = 0$. Condition (b) together with Equation (4) clearly implies that any positive value of $\bar{c}$ results in $\bar{b}_i < 0$. Thus $\mathcal{B}$ is once again unsplittable.

For the converse we will need to show that if either $\sum t_i$ is even or both (a) and (b) are false, then $\mathcal{B}$ is splittable. First we claim that the choice $\bar{c} = \lfloor pc \rceil$ always ensures that $0 \leq \bar{b}_i \leq b_i$. To see this note first that this choice ensures $|pc - \bar{c}| \leq \frac{1}{2}$. Moreover we can always assume $|\rho_i| < \frac{3}{2}$, since otherwise all $|\epsilon_j| = \frac{1}{2}$ and we would be able to change the rounding of the targets t_i (even while preserving the parity of $\sum t_i$). Thus Equation (4) implies that $\bar{b}_i = pb_i + E$ where $E < 1$, and we can therefore argue as in the proof of Lemma 3.1 to complete the claim.

Of course we cannot necessarily choose $\bar{c} = \lfloor pc \rceil$, since this may not have the same parity as $\sum t_j$. Thus we also claim that one of the two choices $\bar{c}_- := \lfloor pc \rceil - 1$ or $\bar{c}_+ := \lfloor pc \rceil + 1$ results in $0 \leq \bar{b}_i \leq b_i$. Indeed the two choices result in values of $pc - \bar{c}_-$ and $pc - \bar{c}_+$ that differ by 2 and have absolute value $\leq \frac{3}{2}$. Meanwhile the ρ_i lie in some interval of length ≤ 2 which is contained in $(-\frac{3}{2}, \frac{3}{2})$. It is an easy exercise to conclude that either all $|pc - \bar{c}_- + \rho_i| < 2$ or all $|pc - \bar{c}_+ + \rho_i| < 2$. Thus one of two choices $\bar{c} = \bar{c}_-$ or $\bar{c}_+$ gives values $\bar{b}_i = pb_i + E$ where $E < 1$, and we are again done as in the previous claim.

Now assume that $\sum t_i$ is even. Then if $\lfloor pc \rceil$ is even we have already argued $\bar{c} = \lfloor pc \rceil$ leads to a solution. And if $\lfloor pc \rceil$ is odd then we always have $0 \leq \lfloor pc \rceil \pm 1 \leq c$ (here we are using $p \leq \frac{1}{2}$). Thus one of the choices $\bar{c} = \lfloor pc \rceil \pm 1$ leads to a solution as well.

Finally assume both (a) and (b) are false. Since (a) is false and $p \leq \frac{1}{2}$ we have that $\lfloor pc \rceil + 1 \leq c$. On the other hand since (b) is false we either have (i) $pc \geq \frac{1}{2}$, or else (ii) $pc < \frac{1}{2}$ and all $pb_i + \frac{1}{2}(pc - 1 + \rho_i) \geq 0$. In case (i) we have $0 \leq \lfloor pc \rceil \pm 1 \leq c$, which we have previously shown implies $\mathcal{B}$ is splittable. In case (ii) we have $\lfloor pc \rceil + 1 = 1$ and moreover that the choice $\bar{c} = 1$ leads to a valid solution for all $\bar{b}_i$. This concludes the proof. □

In the lemma, if any of the $p|B_i|$ is a half-integer, then $\mathcal{B}$ is splittable. Indeed, in this case we can select the set target t_i to make $\sum t_i$ even. We also record here that if case (b) of the lemma holds, then there is in fact a unique i such that $2pb_i + pc - 1 + \rho_i < 0$, and moreover we must have that $\rho_i < -1$, that $\epsilon_i > 0$, and that the other two $\epsilon_j < 0$.

In the next result we consider the case when a collection $\mathcal{B}$ of three sets has elements of multiplicity 1.

Lemma 3.6. *Assume that $p \leq \frac{1}{2}$, and let $\mathcal{B} = \{B_1, B_2, B_3\}$ be a given collection with a_i, b_i, c as in Figure 2. Then provided at least one of the a_i is sufficiently large, $\mathcal{B}$ is p-splittable.*

Proof. First let $\mathcal{B}^{(0)}$ be the collection obtained from $\mathcal{B}$ by removing all elements of multiplicity 1. For the rest of the proof, let b_i, c, t_i, ρ_i be as in the previous lemma, *for the collection $\mathcal{B}^{(0)}$*.

Suppose first that $\mathcal{B}^{(0)}$ is p-splittable. Then $\mathcal{B}$ is splittable too; in fact we claim that any splittable collection remains splittable after adding elements of multiplicity 1. To see this, it suffices to show it when we add just one element a of multiplicity 1 to some set B_i. Now if adding a raises the value of t_i by 1, then we include a in the splitter; otherwise we would exclude a from the splitter.

Next suppose that $\mathcal{B}^{(0)}$ is p-unsplittable. Then by Equation (3) we can "split" the collection $\mathcal{B}^{(0)}$ by setting $\bar{c} = 0$ and $\bar{b}_i = \frac{1}{2}\left(-t_i + \sum_{j \neq i} t_j\right)$. However these choices of $\bar{b}_i$ will be half-integers, and need not satisfy $0 \leq \bar{b}_i \leq b_i$.

We claim that for all i we have $-\frac{1}{2} \leq \bar{b}_i \leq b_i + \frac{1}{2}$, and that there is at most one i such that $\bar{b}_i = -\frac{1}{2}$ or $\bar{b}_i = b_i + \frac{1}{2}$. The first part of the claim follows directly from Equation (4), together with $\bar{c} = 0$ and $pc < \frac{1}{2}$.

For the second part of the claim, the only plausible contrary case is one where two of the $\bar{b}_i$ have opposite error, say $\bar{b}_2 = -\frac{1}{2}$ and $\bar{b}_3 = b_3 + \frac{1}{2}$. We now show this implies that $b_3 = 0$. Indeed, Equation (4) for $\bar{b}_2$ says $-\frac{1}{2} = pb_2 + \frac{1}{2}(pc + \rho_2)$ and this implies $\rho_2 < -1$. Since all ρ_i lie within an interval of length 2, it follows that $\rho_3 < 1$. Then Equation (4) for $\bar{b}_3$ says $b_3 + \frac{1}{2} = pb_3 + \frac{1}{2}(pc + \rho_i)$. Now $p \leq \frac{1}{2}$, $pc < \frac{1}{2}$, and $\rho_3 < 1$ all together imply $b_3 = 0$.

Using $b_3 = 0$, we obtain in particular that $t_3 \geq t_2$. On the other hand Equation (3) for $\bar{b}_2$ says $-1 = t_1 - t_2 + t_3$, and this implies $t_2 > t_3$. This is a contradiction, and completes the proof of the claim.

Now we can finish the proof as follows. Let $\mathcal{B}^{(1)}$ be the collection obtained from $\mathcal{B}$ by removing just the elements of B_2 and B_3 of multiplicity 1. That is, $\mathcal{B}^{(1)}$ is obtained by zeroing out a_2 and a_3. We will show that if a_1 is sufficiently large, then $\mathcal{B}^{(1)}$ is splittable.

For this, we will use the notation $t_1^{(1)}$ for the target value of B_1 as it would be defined for $\mathcal{B}^{(1)}$ (or equivalently for $\mathcal{B}$). Thus in particular $t_1^{(1)} \geq t_1$. In the next paragraph we will only have need of small values of a_1, so that we need only consider the cases $t_1^{(1)} = t_1$ and $t_1^{(1)} = t_1 + 1$.

Our above claim implies that at least one of the two triples $\{\bar{b}_1 - \frac{1}{2}, \bar{b}_2 + \frac{1}{2}, \bar{b}_3 + \frac{1}{2}\}$ or $\{\bar{b}_1 + \frac{1}{2}, \bar{b}_2 - \frac{1}{2}, \bar{b}_3 - \frac{1}{2}\}$ lies within the desired bounds $[0, b_1], [0, b_2], [0, b_3]$. In the first case if a_1 is large enough that $t_1^{(1)} = t_1 + 1$ then the triple is a valid splitting. In the second case if $a_1 = 1$ and $t_1^{(1)} = t_1$ then the triple extends to a valide splitting by selecting the single element of a_1. And if $a_1 \geq 2$ and $t_1^{(1)} = t_1 + 1$ then the triple extends to a valid splitting by selecting two elements from a_1.

Thus we have shown in each case that there exists a value of a_1 that results in $\mathcal{B}^{(1)}$ being splittable. By our claim from the second paragraph, any larger value of a_1 will also result

in $\mathcal{B}^{(1)}$ being splittable. Again using the claim from the second paragraph, this always implies $\mathcal{B}$ is splittable. □

The above lemma may seem natural, since intuitively the presence of elements of multiplicity 1 makes it easier to find a splitter. However the analogous result is false for collections of four or more sets. Indeed an unsplittable collection of Type 0 of Appendix A may have a Venn region of multiplicity 1 of arbitrary size.

We claim that the argument of the previous lemma is optimal. That is, in the case when $\mathcal{B}^{(0)}$ is unsplittable, $\mathcal{B}$ will be splittable if and only at least one of the a_i as at least as large as we needed in our proof. We omit the case-by-case justification of this fact.

Thus the results of this section provide a simple procedure to decide whether a given configuration $\mathcal{B}$ is p-splittable. First one can replace p with $1-p$ if necessary to assume that $p \leq \frac{1}{2}$. Next given $\mathcal{B}$ one can obtain $\mathcal{B}^{(0)}$ as in Lemma 3.6, and use Lemma 3.5 to decide whether $\mathcal{B}^{(0)}$ is splittable. If it is, then $\mathcal{B}$ is splittable, and if not, then one can use the details of the proof of Lemma 3.6 to determine whether any a_i is large enough to make $\mathcal{B}$ is splittable.

§4. $\frac{1}{2}$-splittability criteria and characterizations

In the previous section, we examined p-splittability criteria for arbitrary p. In this section we specialize to the important case $p = \frac{1}{2}$. After providing another very general lemma, we use it to give a complete characterization of splittability for collections of at most four sets. Throughout this section, the term splittability will always refer to $\frac{1}{2}$-splittability.

The following result, while quite simple, is useful for converting our understanding of collections with few elements into more general theorems.

Lemma 4.1 (Reduction Lemma)**.** *Let $\mathcal{B}$ be a given collection, and let $\mathcal{B}'$ be a collection obtained from $\mathcal{B}$ by adding an even number of elements to any of its Venn regions. Then if $\mathcal{B}$ is splittable, so is $\mathcal{B}'$.*

Proof. If S is a splitter for $\mathcal{B}$, then we can construct a splitter S' for $\mathcal{B}'$ as follows. Begin by putting all the elements of S into S'. Then for each Venn region R of $\mathcal{B}$ and corresponding Venn region R' of $\mathcal{B}'$, put half of the elements of $R' \setminus R$ into S'. It is easy to see that S' is a splitter for $\mathcal{B}'$. □

Before stating our characterization of splittability for configurations with four sets, we review the known characterization of splittability for configurations with three or fewer sets.

Proposition 4.2. *Any collection of one or two sets is splittable. A collection of three sets is unsplittable if and only if both:*

(a) *every Venn region of multiplicity* 2 *contains an odd number of elements; and*
(b) *all other Venn regions are empty.*

The proposition can easily be extracted from the results of the previous section. It is also possible to give a direct proof, as was done in [CCSS16].

Theorem 4.3. *Any family of four sets is unsplittable if and only if it falls into one of the eleven types described in Appendix A.*

Proof. First we claim that each of the types described in Appendix A is unsplittable. The collections of Type 0 are by definition those which have some three-set subcollection that is unsplittable. It is tedious but straightforward to check that all of Types 1–10 are unsplittable; as an example we provide the proof that Type 1 is unsplittable in Proposition A.1.

In order to prove that all families not matching Types 0–10 are in fact splittable, note first that we have used a supercomputer to test splittability for every collection of four sets with all Venn regions of size ≤ 3. The code and its output are available in [REU16].

Now suppose that $\mathcal{B}$ is a collection of four sets which is not of any of Types 0–10. We wish to show that $\mathcal{B}$ is splittable. Let $\mathcal{B}^{(2)}$ be a collection obtained from $\mathcal{B}$ emptying each Venn region R that is even in $\mathcal{B}$, and leaving just 1 element in each Venn region R that is odd in $\mathcal{B}$. In other words, $\mathcal{B}^{(2)}$ is obtained by taking each Venn region "modulo 2".

If $\mathcal{B}^{(2)}$ is also not of any of the eleven types, then since its regions all have size ≤ 3 our software has checked that $\mathcal{B}^{(2)}$ is splittable. By the Reduction lemma (Lemma 4.1), $\mathcal{B}$ is also splittable, and we are done in this case.

On the other hand, suppose that $\mathcal{B}^{(2)}$ is of one of the eleven types, say type T. Then since $\mathcal{B}$ is not of type T, there must exist a Venn region R, labelled 0 (empty) or 1 or 0/1, such that the size of R in $\mathcal{B}^{(2)}$ is strictly less than the size of R in $\mathcal{B}$. Now let $\mathcal{B}'$ be the configuration obtained from $\mathcal{B}^{(2)}$ by adding 2 elements to R. Thus $\mathcal{B}'$ is not among the Types 0–10. And since the regions of $\mathcal{B}'$ still have size ≤ 3, our software has checked that $\mathcal{B}'$ is splittable. It again follows from Lemma 4.1 that $\mathcal{B}$ is splittable, which concludes the proof. □

Looking at the unsplittable types of collections of four sets in Appendix A, one might surmise that unsplittable configurations should have many Venn regions with few or zero elements. We next provide a result which says that if a collection has sufficiently many Venn regions which are sufficiently large, then it must be splittable.

Theorem 4.4. *Let D be an integer bound on the discrepancy of collections of n sets. Suppose that $\mathcal{B}$ is a collection of n sets such that each Venn region of multiplicity 1 contains at least $D-1$ elements. Then $\mathcal{B}$ is splittable.*

Proof. Let $\mathcal{B} = \{B_1, \ldots, B_n\}$ be such a collection, and let R_i denote the Venn region of multiplicity 1 contained in B_i. Then in $\mathcal{B}$, each region R_i has at least $D-1$ elements, so we may let $\mathcal{B}^{(0)} = \{B_1^{(0)}, \ldots, B_n^{(0)}\}$ be the collection obtained from $\mathcal{B}$ by deleting $D-1$ elements from each of the Venn regions R_i.

Since $\mathrm{disc}(\mathcal{B}^{(0)}) \leq D$, we can find a set $S^{(0)}$ such that for all i we have

$$-D \leq |B_i^{(0)} \cap S^{(0)}| - |B_i^{(0)} \setminus S^{(0)}| \leq D$$

Now for each i we restore the D' deleted elements of the Venn region R_i. As we do so, we build a set S by beginning with $S^{(0)}$, and then placing *some* of the $D-1$ restored elements into S and the rest into S^c. It is easy to do so in such a way that

$$-1 \leq |B_i \cap S| - |B_i \setminus S| \leq 1$$

and as a result S splits $\mathcal{B}$. □

Of course in the above result, D can be taken to be the ceiling of Spencer's bound $K\sqrt{n}$ discussed in the introduction.

§4.1. The prevalence of splittability. In this subsection we address several questions about how commonly splittable and unsplittable collections occur. Our results for small collections of sets indicate that unsplittability is very rare. It is natural to ask whether this remains true for collections with a larger number of sets.

We first consider the prevalence of splittability when the number of elements is fixed and the number of sets is large. In the following result, we let $f(n,k)$ denote the fraction of all n-set collections on k elements which are splittable.

Proposition 4.5. *Suppose that $k \geq 3$. Then $f(n,k) \to 0$ as $n \to \infty$.*

Intuitively, the proposition holds because there exists a three-set configuration with three elements that is unsplittable, namely $\mathcal{B}_0 = \{\{1,2\},\{1,3\},\{2,3\}\}$. As the number of sets in the collection increases, the probability that its restriction to the points $1,2,3$ contains the collection $\mathcal{B}_0$ becomes highly likely.

On the other hand, if we fix the number n of sets and let $k \to \infty$ then the argument of the next result implies that $f(n,k)$ converges to 1. In fact, the same holds even if we let $n, k \to \infty$ with k growing fast enough with respect to n.

Theorem 4.6. *Suppose that $k = k(n)$ lies in $\omega(2^n n)$, that is, k grows asymptotically strictly faster than $2^n n$. Then $f(n,k) \to 1$ as $n \to \infty$.*

Proof. Let $D = K\sqrt{n}$ where K is the constant of Theorem 4.4. Our strategy is to show that if k is as large as in our hypothesis, then it is unlikely that any of the multiplicity 1 regions will contain fewer than D elements. Thus by Theorem 4.4 it is likely that a given configuration will be splittable.

For this note that if a single element is randomly assigned to each of n sets, then the probability that the element will lie in any given Venn region is $q = \frac{1}{2^n}$. Next assign k elements randomly and independently to the sets, and let the random variable X denote the number elements of a fixed Venn region of multiplicity 1. By the basic properties of the binomial distribution, the expected value of X is $\mu = kq = \frac{k}{2^n}$ and the standard deviation of X is $\sigma = \sqrt{kq(1-q)} = \sqrt{k\frac{1}{2^n}(1-\frac{1}{2^n})}$.

We now wish to bound from above $\Pr[X < D]$. Letting $t = \frac{\mu - K\sqrt{n}}{\sigma}$, we have that $\Pr[X < D] \leq Pr[|X - \mu| \geq t\sigma]$. Chebychev's inequality now gives

$$\Pr[X < D] \leq \frac{1}{t^2}$$

Substituting the expressions for t, μ, σ and simplifying we obtain

$$\begin{aligned}\Pr[X < D] &\leq \frac{\frac{k}{2^n}(1-\frac{1}{2^n})}{(\frac{k}{2^n} - K\sqrt{n})^2} \\ &\leq \frac{2^n}{k - 2K\sqrt{n}}\end{aligned}$$

Our hypothesis about the growth of k implies that the latter quantity is $o(1/n)$. Finally the probability that any of the n regions of multiplicity 1 has fewer than D elements is bounded by $n \Pr[X < D]$ and is thus $o(1)$, or in other words, converges to 0. □

Appendix A. Unsplittable configurations of four sets

Here we provide a catalogue of the unsplittable configurations of four or fewer sets. Our catalogue consists of labelled four-set Venn diagrams, shown as four-lobed "hearts" with each lobe representing one set. The diagram below shows four of the diagrams; one with each of the four sets shaded.

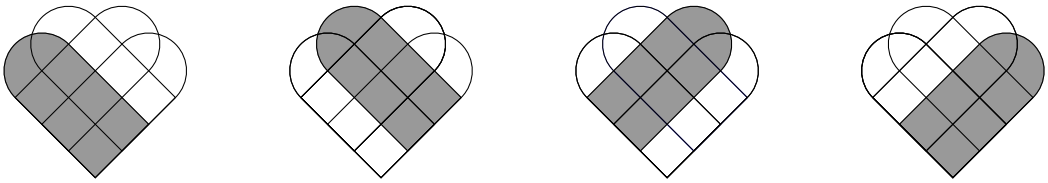

Figure 3 illustrates examples of each type of unsplittable configuration of four sets. In the diagrams we use the following abbreviations: the symbol o denotes an odd number of elements, e denotes an even number of elements, 1 denotes one element, 0/1 denotes zero or one element, x denotes any number of elements, and a blank denotes zero elements. Note that two separate instances of a symbol do *not* necessarily denote the same quantity.

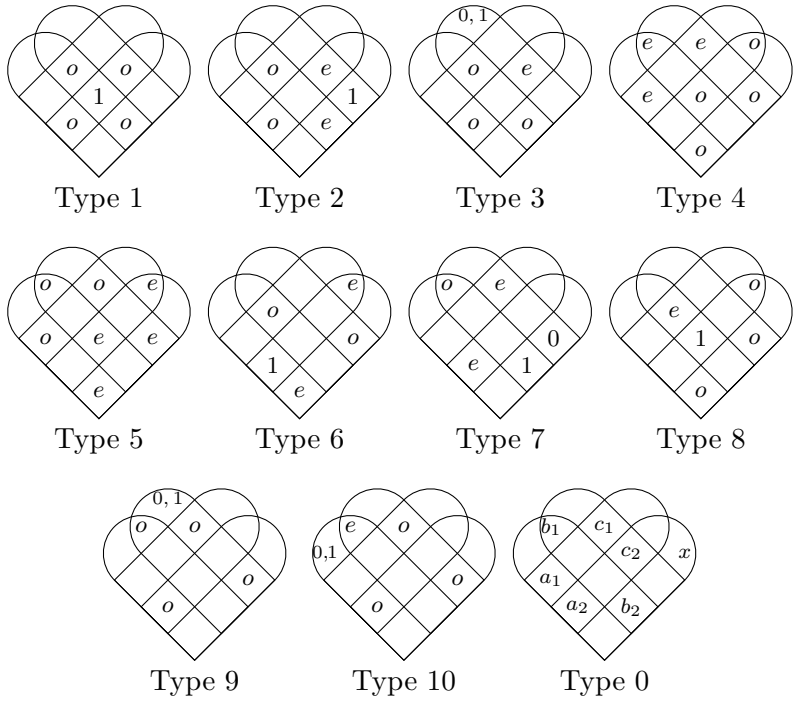

FIGURE 3. Unsplittable 4-set configurations

Some additional remarks about the types are in order. First, each of the types has analogous instances in which the sets $B_1, \ldots, B_4$ are permuted.

In Type 0, we additionally require that $a_1 + a_2$, $b_1 + b_2$, and $c_1 + c_2$ all to be odd numbers. This type covers the cases when some subcollection of three of the sets is unsplittable. In

the instance depicted, the first three sets form an unsplittable configuration. There is some overlap between Type 0 and degenerate instances of other types.

We now briefly address the claim, needed for Theorem 4.3, that configurations lying in any of the eleven types are indeed unsplittable. The unsplittability of Type 0 is handled in Prop 4.2. The unsplittability of Types 4 and 5 follows from Lemma 3.1. The unsplittability of the remaining types all boil down to elementary linear algebra, and as an example let us verify Type 1 is unsplittable.

Proposition A.1. *If $\mathcal{B}$ is a collection of Type 1, then $\mathcal{B}$ is unsplittable.*

Proof. Let o_i denote the number of elements of $\bigcap_{j\neq i} B_j \cap B_i^c$, so that o_i is an odd number. Let $t_i = \frac{1}{2}(1 + \sum_{j\neq i} o_i)$ be the target quantity for $S \cap B_i$. Then finding a splitter S for $\mathcal{B}$ is equivalent to solving the integer system:

$$\begin{bmatrix} 1 & 1 & 1 & & 1 \\ 1 & 1 & & 1 & 1 \\ 1 & & 1 & 1 & 1 \\ & 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ b \end{bmatrix} = \begin{bmatrix} t_1 \\ t_2 \\ t_3 \\ t_4 \end{bmatrix}$$

subject to the constraints that $0 \leq a_i \leq o_i$, and $0 \leq b \leq 1$. Solving this system for a_1 in terms of the right-hand sides and b, we obtain the equation:

$$\begin{aligned} a_1 &= \frac{1}{3}(t_1 + t_2 + t_3 - 2t_4 - b) \\ &= \frac{1}{3}\left(\frac{o_1 + o_2 + o_3 + 1}{2} + \frac{o_1 + o_3 + o_4 + 1}{2} + \frac{o_1 + o_2 + o_4 + 1}{2} - 2\frac{o_2 + o_3 + o_4 + 1}{2} - b\right) \\ &= \frac{1}{2}o_1 + \frac{1}{6}(1 - b) \end{aligned}$$

If $b = 0$ then this implies $a_1 = \frac{1}{6}(3o_1 + 1)$. This is a contradiction since o_1 odd implies that $3o_1 + 1 \not\equiv 0 \pmod 6$. On the other hand if $b = 1$ then $a_1 = \frac{1}{2}o_1$, which is clearly a contradiction from o_1 odd. □

References

[AS00] N. Alon and J. Spencer. *The Probabilistic Method*. Wiley, 2000.

[Ban10] N. Bansal. Constructive algorithms for discrepancy minimization. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, pages 3–10. IEEE Computer Society, 2010.

[BF81] József Beck and Tibor Fiala. "Integer-making" theorems. *Discrete Appl. Math.*, 3(1):1–8, 1981.

[BH97] Debe Bednarchak and Martin Helm. A note on the beck-fiala theorem. *Combinatorica*, 17(1):147–149, 1997.

[Buk13] B. Bukh. An improvement of the beck-fiala theorem. *ArXiv e-prints*, June 2013.

[CCSS16] D. Condon, S. Coskey, L. Serafin, and C. Stockdale. On generalizations of separating and splitting families. *Electronic Journal of Combinatorics*, 23(3):#P3.36, 2016.

[DPV06] S. Dasgupta, C. Papadimitrious, and U. Vazirani. *Algorithms*. McGraw Hill, 2006.

[EH] John Elton and Theodore Hill. Ham sandwich with mayo: A stronger conclusion to the classical ham sandwich theorem.

[Hel99] Martin Helm. On the beck-fiala theorem. *Discrete Mathematics*, 207(1):73 – 87, 1999.

[Kar72] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*, pages 85–103. Plenum, New York, 1972.

[LM12] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. *CoRR*, abs/1203.5747, 2012.

[Lov73] L. Lovász. Coverings and coloring of hypergraphs. In *Proceedings of the Fourth Southeastern Conference on Combinatorics, Graph Theory, and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1973)*, pages 3–12. Utilitas Math., Winnipeg, Man., 1973.

[REU16] 2016 REU. Split checker public repository, 2016.

[Spe85] J. Spencer. Six standard deviations suffice. *Transactions of the AMS*, 289(2):679–706, 1985.

Peter Bernstein, Tufts University, Medford, MA 02155
E-mail address: peter.bernstein@tufts.edu

Cashous Bortner, University of Nebraska, Lincoln, NE 68588
E-mail address: cashous.bortner@huskers.unl.edu

Samuel Coskey, Boise State University, Boise, ID 83702
E-mail address: scoskey@nylogic.org

Shuni Li, Macalester College, Saint Paul, MN 55105
E-mail address: sli@macalester.edu

Connor Simpson, Cornell University, Ithaca, NY 14853
E-mail address: cgs93@cornell.edu